

Gravimetric Localization on the Surface of Small Bodies

Benjamin Hockman
 Dept. of Mechanical Engineering
 Stanford University
 Stanford, CA, 94305
 302-379-6645
 bhockman@stanford.edu

Robert G. Reid
 Jet Propulsion Laboratory
 California Institute of Technology
 Pasadena, CA, 91109
 818-393-5946
 rgreid@jpl.nasa.gov

Issa. A. D. Nesnas
 Jet Propulsion Laboratory
 California Institute of Technology
 Pasadena, CA, 91109
 818-393-5007
 nesnas@jpl.nasa.gov

Marco Pavone
 Dept. of Aeronautics & Astronautics
 Stanford University
 Stanford, CA, 94305
 650-723-4432
 pavone@stanford.edu

Abstract—The localization of landers on the surface of small bodies has traditionally relied on observations from a mothership (e.g. Rosetta’s Philae lander and Hayabusa 2’s MASCOT and MINERVA landers). However, when line-of-sight with the mothership is not always available, or for surface rovers that travel large distances, alternative mothership-independent localization techniques may be required. On-board vision-based techniques have demonstrated effective localization in terrestrial applications as well as for Mars rovers, but may be unreliable on small bodies where rovers must contend with fast-moving shadows, difficulties observing absolute scale, and issues acquiring images such as dust, sun blinding, tumbling and low albedo. We investigate the feasibility of an entirely new localization approach based on surface *gravimetry*, where a rover can constrain its location on the surface by precisely measuring the local gravity vector. This mothership-independent localization technique is well-suited to a class of hybrid rovers that can bounce and tumble over the surface of small bodies; it is insensitive to surface illumination, and even works at night. We develop a Bayesian framework for computing localization “likelihood maps” from gravimetry (and gradiometry) data, accounting for all sensor and model uncertainties. We then propose a method for deriving landing distributions of a bouncing rover from simulation data to serve as a prior for the localization estimate. Finally, we conduct a case study on the Philae lander, where we show how this approach could have helped reject localization hypotheses and significantly narrow the areas searched for the Philae lander.

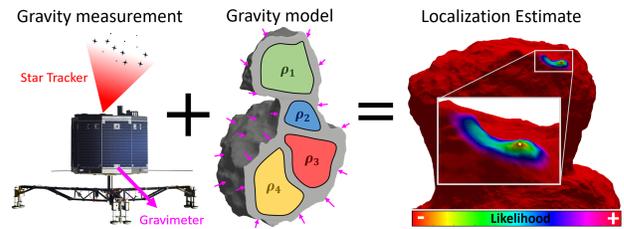


Figure 1. Gravimetric localization concept: a small body lander or rover equipped with (1) a star tracker to measure its celestial orientation, (2) a gravimeter to measure the local gravity vector, and (3) a model of the gravity field can infer its location on the surface.

TABLE OF CONTENTS

1. INTRODUCTION.....	1
2. LOCALIZATION APPROACHES	2
3. GRAVIMETRIC LOCALIZATION	3
4. MEASUREMENT MODEL	3
5. PROBABILISTIC DYNAMICS MODEL	7
6. CASE STUDY: THE PHILAE LANDER	9
7. CONCLUSIONS.....	10
REFERENCES	10
BIOGRAPHY	11

1. INTRODUCTION

Small solar system bodies, such as asteroids, comets, and irregular moons are important targets for planetary exploration for several reasons. First and foremost, small bodies have un-

dergone little weathering since the early formation of the solar system, and therefore could provide unique *science data* to help understand the evolution of the Solar System and origins of life [1]. Second, impact mitigation strategies for *planetary defense* require a better understanding of the structural and compositional properties of Near-Earth Asteroids (NEAs) [2]. Third, many small bodies contain valuable *resources* that could be used in-situ (e.g., to generate propellant or as construction materials) or even for return to Earth [3]. Finally, their low ΔV requirements makes small bodies promising targets for future *human exploration*, potentially as stepping stones to Mars and beyond [4]. Due to this breadth of exploration interests and their relative accessibility, it is not surprising that there have been an increasing number of missions to small bodies from space agencies worldwide.

Launched in 1996, NASA’s Near Earth Asteroid Rendezvous mission (NEAR) was the first to provide a detailed characterization of an asteroid (Eros) and land on its surface [5]. In 1999, NASA’s Stardust mission provided the first sample return from a small body, as it collected dust grains from the coma of Comet Wild 2 [6]. From 2003-2005, a wave of new missions took flight. In 2003, JAXA launched Hayabusa—a spacecraft that characterized the surface of Asteroid 25143 Itokawa, attempted to deploy a small lander (which failed to reach the surface), and returned fine grains of regolith for laboratory analysis [7]. ESA’s Rosetta spacecraft, launched in 2004, began its ten-year voyage to Comet 67P/C-G, where it performed detailed remote characterization of the cometary nucleus and deployed the Philae lander to the surface [8]. In 2005, NASA’s Deep Impact spacecraft was sent to investigate the sub-surface composition of Comet Tempel 1 with a surface impactor [9]. Ongoing missions include NASA’s Dawn

spacecraft, studying two protoplanets in the asteroid belt: Vesta and Ceres [10], JAXA’s Hayabusa 2 spacecraft, which is en route to Asteroid 162173 Ryugu, will deploy four surface landers [7], and NASA’s OSIRIS-REx mission, which is en route to rendezvous with the carbonaceous asteroid Bennu in 2018 and return a sample in 2023 [11]. Planned future missions include NASA’s two 2017 Discovery selections, Lucy [12] and Psyche [13], to explore six Jupiter Trojan asteroids and a metallic asteroid, respectively, and JAXA’s Martian Moons eXploration mission (MMX), which will perform close-proximity characterization and sample return from Mars’ moons Phobos and Deimos [14].

Remote observations of small bodies from orbiting spacecraft can help to constrain some chemical and bulk physical properties, but measurements of composition (e.g. for origin science), physical properties of the regolith, and interior structure require direct contact with the surface at multiple locations for extended periods of time [15]. Accordingly, low-cost landers and rovers could offer new scientific perspectives and, as such, they have been included as secondary payloads on many missions to small bodies.

In addition to deploying a surface impactor, camera, and five target markers, Hayabusa 2—which is scheduled to rendezvous with Asteroid 162173 Ryugu in July 2018—will deploy four surface landers: a 10 kg MASCOT lander developed by DLR and three 500 g MINERVA landers, developed by JAXA. Each lander/rover is equipped with an internal momentum device for self-righting and to perform small hops, albeit with minimal control. Other research groups are developing hopping platforms for more precise mobility, including the Highland Terrain Hopper [16] and Hedgehog rover [17].

However, a key challenge for landers and rovers on small bodies is the reduced gravity environment, or “microgravity.” The failed gravitational capture of the MINERVA lander deployed from Hayabusa and the inadvertent bouncing of the Philae lander on Comet 67P/C-G highlight the difficulties of landing on and remaining attached to small bodies. Indeed, even the Philae lander, which was equipped with three devices designed to attach itself to the surface (harpoon, drills and thrusters) and touched down at only 1 m/s, still bounced over 1 km off target, landing sideways in the shadow of a large rock/cliff.

Whether it’s a hopping rover designed to passively bounce and tumble over the surface, or the Philae lander’s failed anchoring attempt, the stochasticity of surface dynamics presents many challenges for *localization and navigation*. The dynamics of Mars rovers are slow and predictable, so they can reliably navigate using wheel- and vision-based odometry. Bouncing and tumbling rovers, however, must contend with spinning cameras, stark differences in day-to-night surface illumination, and a large distribution of possible settling locations. In some cases, a mothership may be able to aid localization by imaging the landing site (as is planned for Hayabusa 2), but in general, a mothership may have observational constraints such as line-of-sight occlusions from orbit or simply coverage limitations due to the camera’s field-of-view. Philae relied on Rosetta to image the landing sight and confirm its location, which, due to the failed anchoring attempt, took over a year of searching. Ideally, we would like a mothership-independent localization strategy that is robust, while both sensing and processing is performed on-board the rover.

In this paper, we investigate the feasibility of a novel localiza-

tion approach that leverages *surface gravity signatures* to infer a lander/rover’s location (see Fig. 1). The concept, in principle, is simple: a rover equipped with a star tracker to measure its orientation, a gravimeter to measure the local gravity vector, and a shape and gravity model of the body can constrain the locations at which a measurement must have been taken. This concept is critically enabled by (1) the small sizes and irregular shapes of small bodies which produce highly varying surface gravity fields, and (2) the development of small precision microgravity gravimeters. We show that... results

2. LOCALIZATION APPROACHES

Localization approaches vary depending on how “lost” one is and what sensors are available; they can be divided into either local or global techniques: local techniques estimate pose (position and orientation) relative to previous poses, while global techniques estimate pose relative to a single privileged frame and vary depending on the type and confidence of prior information. In terrestrial applications, the Global Positioning System (GPS) is an example of global localization with strong priors (satellite locations).

Without GPS, interplanetary spacecraft typically rely on ground-based localization techniques using two-way radiometric measurements and navigation camera images. On-board vision-based localization around small bodies was demonstrated by the AutoNav system during various phases of the Deep Space 1, Stardust, Deep Impact and Dawn missions [18]. AutoNav matches visual features extracted from both camera images and synthetic images rendered from a detailed shape model prior that was processed on the ground. Another technique that uses strong priors is Terrain Relative Navigation (TRN), which will be used for entry descent and landing on the Mars 2020 mission [19]. TRN estimates global pose by matching descent camera images to a prior map and fusing inertial measurements to provide high-rate localization, velocity and attitude estimates.

TRN-like techniques are not suitable for global localization on the surface of small bodies, however, because they require a downwards-facing view of the surface and prior maps at resolutions that are typically not available. Surface-based localization techniques related to AutoNav have been demonstrated for Mars rovers, however, where camera images are matched to images synthesized from detailed 3D elevation maps (processed on the ground from high-resolution orbital imagery) [20]. While this approach works for rovers with limited mobility and relatively small prior maps, for hopping rovers that could bounce repeatedly and come to rest anywhere on the surface of a small body, these techniques are unlikely to scale due to prior map size.

In many interplanetary situations, global measurements and strong priors are not available, such that only local, relative, localization can be performed. Visual odometry (VO) is a relative localization techniques that has also been demonstrated on Mars [21]. In VO, visual features (2D representations of 3D landmarks) are tracked between successive camera images and both the rovers relative motion and the 3D position of landmarks are robustly estimated. The relative motion estimates from VO are typically accumulated to record a rovers local path, however uncertainties in the rover’s pose will grow without bound. To reduce pose uncertainties, correspondences can be manually established between landmarks (e.g. rocks) viewed in both on-board cameras and orbital imagery.

While the MINERVA rover never validated its localization approach, in [22], Yoshimitsu et al. describes how the rover was designed to estimate its attitude after awakening using six Sun sensors and integrated gyroscope measurements. Relative velocities were to be estimated during hops using optical flow from surface images. Presumably, the MINERVA rover could have estimated both hopping direction and approximate distance, however it is not clear how it was to localize itself globally.

For collaborative missions to small bodies like Hayabusa/MINERVA and Rosetta/Philae, where a mothership deploys a secondary lander (in addition to providing prior maps) the mothership can assist in localizing the lander on the surface using a combination of radiometrics and camera-based surveys (assisted by reflectors or LED strobes). Both the Hayabusa and Rosetta motherships captured images of their landers after deployment. In the case of Rosetta, after Philae’s failed anchoring and subsequent bounces, these images were pivotal in narrowing down the global search area for the lander once combined with line-of-site radiometric measurements [8]. While camera-based surveys, and other synergistic approaches, are valuable for surface localization, they depend on the mothership’s availability (i.e. mission priorities) and its ability to station keep above the surface of a rotating body.

To reduce the dependence on a primary spacecraft, future missions to the surface of small bodies may rely on vision-based techniques, such as VO [23], for relative motion estimates and star trackers for absolute attitude estimates. Several problems that are unique to small bodies must be addressed, however; for example, absolute, metric scale is relatively unobservable from a hopping rover’s on-board sensors. In [23], So et al. show how a stereo camera pair could help determine scale, however with small baselines (10 cm), scale remains unobservable for hops over 10 m in height. Another problem unique to fast-rotating airless bodies is the large numbers of visual feature outliers caused by incorrectly tracking high-contrast moving shadows; these and other problems make vision-based localization highly-prone to scale drift and accumulated position errors. Global (re)localization to a prior map may be required to reduce errors, or after an extended maneuver in a low-power state. This requires the rover to capture a downwards view of the surface with sufficient overlap to the prior map, however it may not be practical or safe to perform such a large ballistic hop. In the next section, we introduce a mothership-independent localization technique that does not require the rover to perform hops or other motions.

3. GRAVIMETRIC LOCALIZATION

“Gravimetry,” or the measurement of a gravitational field, has a long history in terrestrial geophysics for detecting and mapping sub-surface density distributions. With extremely high sensitivities (roughly $1 \mu\text{Gal}$ or $10^{-9} g$), gravimeters can detect perturbations in the nominal gravity field and thus, infer local sub-surface mass/density anomalies. Moreover, by making a large set of distributed measurements over a particular region, gravity inversion techniques can be used to, for example, map groundwater reservoirs, prospect for petroleum and other minerals, or reconstruct the sub-surface geophysical structure [24]. Extending gravity inversion techniques to small body environments was originally proposed by [25] and enabled by the development of a small ($10 \times 10 \times 10 \text{ cm}$, 1 kg), precise ($1 \mu\text{Gal}$) microgravity gravimeter.

Gravimetric surveys and gravity inversion techniques require

that the sensors’ locations be known precisely. The inverse is also true: if the gravity field is assumed to be known, then the gravity measurements provide information about the sensor’s possible location. However, in practice, both the gravity model and gravimeter location may have some uncertainty. This inherent chicken-and-egg problem of estimating both the rover’s position and gravity model is algorithmically similar to the “simultaneous localization and mapping” (SLAM) problem in robotics—a classical joint-estimation problem with a rich body of literature and practical success in many domains, including self-driving cars, unmanned aerial vehicles, and indeed even Mars rovers. The full gravimetric SLAM problem will be left for future work. In this paper, we focus on the more fundamental problem of deriving a probabilistic localization likelihood map given some prior *estimate* of the gravity field and its *uncertainty*.

We cast the localization problem as a simple Bayesian filter,

$$p(\mathbf{x}|\mathcal{o}) = \frac{p(\mathcal{o}|\mathbf{x})p(\mathbf{x})}{\int p(\mathcal{o}|\mathbf{x}')p(\mathbf{x}')d\mathbf{x}'} = \eta p(\mathcal{o}|\mathbf{x})p(\mathbf{x}), \quad (1)$$

where the probability of being at state x given observations \mathcal{o} depends on the measurement model, $p(\mathcal{o}|x)$, and prior, $p(x)$. In general, \mathcal{o} may include information from both onboard sensors such as cameras and solar panels, and also from a mothership (e.g. derived from Doppler/ranging, or surface imaging). In Sect. 4, we focus only on the measurement model of a gravimeter/star tracker, which may be used exclusively, or as a supplement to other sensors. In Sect. 5, we discuss an approach for constructing a probabilistic prior, $p(x_t|x_{t-1}, v_{t-1})$, based on Monte Carlo simulations from some previously known position and velocity—an equally important component for computing good localization estimates. Finally in Sect. 6 we combine the measurement model and prior in a notional case study of the Philae lander, showing that even with modest uncertainties, a gravimeter can be used to generate high-quality localization estimates.

4. MEASUREMENT MODEL

A gravimeter at rest on the surface of a rotating body measures the superposition of gravity and inertial acceleration. The relative scale of inertial acceleration at Earth’s equator is only 0.3% of gravity and is ignored in most applications. However, the size and rotation rates of small bodies vary widely—some even have unstable equatorial regions (e.g. KW4 Alpha). Thus, it is critical to understand both the gravity field *and* rotation state of the body. We denote the total coordinate acceleration at point x in the body frame as

$$\mathbf{a}(x) = \mathbf{g}(x) - N\mathbf{a}^x. \quad (2)$$

Figure 2 depicts the geometry of the problem, whereby a lander (L) is at rest on body B , which is rotating with respect to the celestial (inertial) reference frame. Right-handed bases \hat{l} , \hat{b} , and \hat{c} are attached, respectively.

With a gravimeter alone, our measurement is simply $\mathcal{o} := \{\bar{\mathbf{a}}\}$, where the bar denotes measured quantities. The measurement model is given by,

$$p(\bar{\mathbf{a}}|\mathbf{x}) = \int_{\mathbf{a}'} p(\bar{\mathbf{a}}|\mathbf{a}(x) = \mathbf{a}') p(\mathbf{a}(x) = \mathbf{a}') d\mathbf{a}', \quad (3)$$

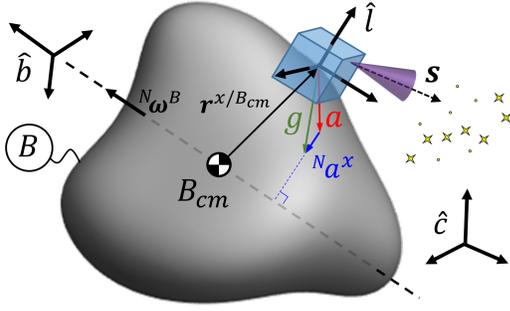


Figure 2. Schematic of problem geometry. A lander (L) is at rest on rotating body (B) and is equipped with a gravimeter to measure the local acceleration, \mathbf{a} , and its celestial orientation (R_c^l) with a star tracker pointed along \mathbf{s} .

where the first term in the integral represents the *sensor model* and captures uncertainties in the gravimeter itself and errors associated with orienting the measured vector in the body frame (i.e. body pose and star tracker errors). The second term represents the *acceleration model* and captures uncertainties in both the gravity field and spin state of the body.

A. Sensor Model

The sensor uncertainty model, $p(\bar{\mathbf{a}}|\mathbf{a}(x) = \mathbf{a}')$, captures the probability distribution of the acceleration estimate, $\bar{\mathbf{a}}$, given the true acceleration, \mathbf{a}' . However, because \mathbf{a} needs to be expressed in \hat{b} and the gravimeter is inherently a measurement in the lander frame, \hat{l} , $\bar{\mathbf{a}}_l$ must be rotated first into the celestial frame via the star tracker (R_c^l), and then into the body frame given the body's celestial pose, R_b^c :

$$\bar{\mathbf{a}}_b = R_b^c R_c^l \bar{\mathbf{a}}_l, \quad (4)$$

where subscripts b and l explicitly denote the frame representation of \mathbf{a} . In this form, $\bar{\mathbf{a}}_b$ has three potential sources of uncertainty: (1) the error model of the gravimeter itself, $p(\bar{\mathbf{a}}_l|\mathbf{a}'_l)$, (2) the rotation uncertainty of the star tracker, and (3) the body pose uncertainty

Gravimeter Error Model

The Vector Gravimeter/Accelerometer (VEGA) instrument developed by Gedex [25] is a 1U microgravity gravimeter that consists of two gimbal-mounted accelerometers to eliminate time-varying bias. When integrated over several minutes in microgravity (below roughly $1\text{ }m\text{g}$), the RMS error is expected to be $1\text{-}10\text{ }n\text{g}^2$. More specifically, VEGA's "error power spectral density" is measured to be about $10^{-14}\text{ }g^2/\text{Hz}$, so for an integration time of say, 10 minutes ($1/600\text{ Hz}$), Parseval's Theorem states that $RMS = \sqrt{10^{-14}/600} = 4.1\text{ }n\text{g}$. This is about three orders of magnitude less than the surface gravity on small asteroids (e.g. for Itokawa, $|g| \approx 2 - 9\mu\text{g}$) and over five orders of magnitude on larger bodies such as Phobos ($|g| \approx 600\mu\text{g}$). If more precision is needed, the sensor can simply be integrated for longer. Since the accelerometers are spherically gimballed, VEGA's uncertainty can be well modeled as a Gaussian with spherical covariance, $\bar{\mathbf{a}}_l \sim \mathcal{N}(\mathbf{a}'_l, \Sigma_{\bar{\mathbf{a}}_l} = \sigma_a^2 I)$.

Star Tracker Error Model

The rotation matrix between the lander and celestial frames is provided by the star tracker, and can be parametrized by Euler

angles, ε_x , ε_y , and ε_z , for rotations about the l_x , l_y , and l_z axes respectively,

$$R_c^l = R_z(\varepsilon_z)R_y(\varepsilon_y)R_x(\varepsilon_x). \quad (5)$$

The accuracy of star trackers is often an order of magnitude better about the pointing axes (i.e. cross-boresight) than it is about the roll axis. For example, the NST-1 nano star tracker has a 3σ pointing accuracy of 7 arcsec but a roll accuracy of 70 arcsec. Thus, it is convenient to define \hat{l} such that one of its principle axes are aligned with the boresight axis, say $\mathbf{s} = \hat{l}_z$. The rotation covariance can then assume a diagonal form, $\Sigma_\varepsilon = \text{diag}([\sigma_{\varepsilon_c}^2, \sigma_{\varepsilon_c}^2, \sigma_{\varepsilon_b}^2])$, where $\sigma_{\varepsilon_b}^2$ and $\sigma_{\varepsilon_c}^2$ are the boresight and cross-boresight variances, respectively.

Body Pose Error Model

The rotation between the celestial and body frames, R_b^c , may also have some uncertainty. Again, we can use Euler angles ξ_x , ξ_y , and ξ_z to parametrize this rotation,

$$R_b^c = R_z(\xi_z)R_y(\xi_y)R_x(\xi_x), \quad (6)$$

where, by convention, \hat{b}_z is aligned with the body's spin axis, and \hat{b}_x is typically aligned with its minor axis. Again, this also diagonalizes the pose covariance Matrix, $\Sigma_\xi = \text{diag}([\sigma_{\xi_a}^2, \sigma_{\xi_a}^2, \sigma_{\xi_t}^2])$, where $\sigma_{\xi_a}^2$ represents the variance in the direction of the spin axis (ξ_x , ξ_y), and $\sigma_{\xi_t}^2$ is the variance associated with the rotation phase about that axis (ξ_z). In practice, through extensive remote observation before lander deployment, a mothership can estimate the body's rotation and pose quite accurately. Moreover, a stationary lander can use its star tracker to take successive images over time (perhaps one full period) to further refine the body's spin and pose estimate.

Note that it is not necessary that R_b^c and R_c^l be computed at the exact time of the gravimetry measurement, but only that they be estimated simultaneously (i.e. at the timestamp of the star tracker image).

Total Sensor Uncertainty

Given uncertainty models for the gravimeter, star tracker, and body pose, the uncertainty of $\bar{\mathbf{a}}_b$ from Eq. (4) is given by,

$$\Sigma_{\bar{\mathbf{a}}_b} = \mathbf{J}_{\bar{\mathbf{a}}_l} \Sigma_{\bar{\mathbf{a}}_l} \mathbf{J}_{\bar{\mathbf{a}}_l}^T + \mathbf{J}_\varepsilon \Sigma_\varepsilon \mathbf{J}_\varepsilon^T + \mathbf{J}_\xi \Sigma_\xi \mathbf{J}_\xi^T, \quad (7)$$

where $\mathbf{J}_{\bar{\mathbf{a}}_l} = \partial f / \partial \bar{\mathbf{a}}_l$, $\mathbf{J}_\varepsilon = \partial f / \partial \varepsilon$, $\mathbf{J}_\xi = \partial f / \partial \xi$, and $f = R_b^c R_c^l \bar{\mathbf{a}}_l$. Since rotation matrices R_b^c and R_c^l can be decomposed into a series multiplication of single-axis rotations (i.e. Eqs. (5) and (6)), the Jacobians can also be constructed purely through matrix multiplication of simple rotations. For example,

$$\mathbf{J}_\xi = R_c^l \left[R_{\xi_z} R_{\xi_y} \frac{\partial R_{\xi_x}}{\partial \xi_x} \bar{\mathbf{a}}_l \mid R_{\xi_z} \frac{\partial R_{\xi_y}}{\partial \xi_y} R_{\xi_x} \bar{\mathbf{a}}_l \mid \frac{\partial R_{\xi_z}}{\partial \xi_z} R_{\xi_y} R_{\xi_x} \bar{\mathbf{a}}_l \right],$$

$$R_{\xi_x} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \xi_x & -\sin \xi_x \\ 0 & \sin \xi_x & \cos \xi_x \end{bmatrix}, \quad \frac{\partial R_{\xi_x}}{\partial \xi_x} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & -\sin \xi_x & -\cos \xi_x \\ 0 & \cos \xi_x & -\sin \xi_x \end{bmatrix}.$$

²A flight campaign in LEO is scheduled to verify this precision.

B. Inertial Acceleration Model

The acceleration of a point (x) on a rotating rigid body (B) is

$${}^N\mathbf{a}^x = {}^N\boldsymbol{\alpha}^B \times \mathbf{r}^{x/B_{\text{cm}}} + {}^N\boldsymbol{\omega}^B \times ({}^N\boldsymbol{\omega}^B \times \mathbf{r}^{x/B_{\text{cm}}}), \quad (8)$$

where ${}^N\boldsymbol{\omega}^B$ is the angular velocity of B , ${}^N\boldsymbol{\alpha}^B$ is its angular acceleration, $\mathbf{r}^{x/B_{\text{cm}}}$ is the vector from B 's center of mass, B_{cm} , to x , and left superscript N denotes rates with respect to the inertial reference frame³. Note that although most small bodies are in near-constant stable spin about their major axis and α can be neglected, this may not be the case for some “non-principle-axis rotators” where nutation effects are non-negligible. Thus, for generality, we will keep the angular acceleration term.

All terms in Eq. (8) are measured, or derived from measured quantities, and thus have uncertainty. For this analysis, we will assume that ${}^N\boldsymbol{\alpha}^B$, ${}^N\boldsymbol{\omega}^B$, and $\mathbf{r}^{x/B_{\text{cm}}}$ can be well approximated by Gaussian random variables with means $\mu_\alpha, \mu_\omega, \mu_r$ and covariances $\Sigma_\alpha, \Sigma_\omega, \Sigma_r$, respectively. Furthermore, the position vector is *relative*, i.e. $\mathbf{r}^{x/B_{\text{cm}}} = \mathbf{r}^{x/P} - \mathbf{r}^{P/B_{\text{cm}}}$, where P is an arbitrary fixed point on B . Thus, position uncertainty is a convolution of uncertainty in x (e.g. from errors in the mesh model) and the center of mass: $\Sigma_r = \Sigma_x + \Sigma_{B_{\text{cm}}}$ (convolution of Gaussians is additive). In practice, Σ_α is typically very small, $\Sigma_{B_{\text{cm}}}$ is large only along the $\boldsymbol{\omega}$ axis, and Σ_ω is extremely small along the $\boldsymbol{\omega}$ axis (i.e. precise period estimate) and generally less than 1° orthogonal to $\boldsymbol{\omega}$. The size of Σ_x largely depends on the resolution and accuracy of the surface mesh and is discussed further in the next section.

If we define the cross product in matrix form,

$$\mathbf{a} \times \mathbf{b} = \begin{bmatrix} 0 & -a_z & a_y \\ a_z & 0 & -a_x \\ -a_y & a_x & 0 \end{bmatrix} \begin{bmatrix} b_x \\ b_y \\ b_z \end{bmatrix} = M_a \mathbf{b},$$

where ${}^N\boldsymbol{\alpha}^B$, ${}^N\boldsymbol{\omega}^B$, and $\mathbf{r}^{x/B_{\text{cm}}}$ have similar forms, M_α , M_ω , and M_r , respectively, it can be shown that the acceleration uncertainty from Eq. (8) is,

$$\Sigma_{N\mathbf{a}^x} = M_r \Sigma_\alpha M_r^T + M_\alpha \Sigma_r M_\alpha^T + M_\omega^2 \Sigma_r M_\omega^{2T} + (M_r M_\omega - 2M_\omega M_r) \Sigma_\omega (M_r M_\omega - 2M_\omega M_r)^T. \quad (9)$$

Thus, assuming Gaussian uncertainty models for ${}^N\boldsymbol{\alpha}^B$, ${}^N\boldsymbol{\omega}^B$, and $\mathbf{r}^{x/B_{\text{cm}}}$, the inertial acceleration is also Gaussian with mean given by Eq. (8) and covariance given by Eq. (9).

C. Gravity Model

Finally, and perhaps most importantly, the quality of gravimetric localization estimates critically depend on the accuracy of the gravity model: $p(\mathbf{g}(x) = \mathbf{g}')$. Gravity modeling has a long history in astrodynamics and methods range in complexity and accuracy depending on the application. Fundamentally, the gravitational field of an arbitrary body is a sum of the gravitational contributions of each “mass element” of that body, and is typically defined by its potential field, $U(\mathbf{r}) = G \int_M \frac{1}{|\mathbf{r} - \mathbf{r}'|} dm'$, where \mathbf{r} is the position vector of the field point. While this expression converges to a spherical model— $U(\mathbf{r}) = GM/|\mathbf{r} - \mathbf{r}_{\text{cm}}|$ —for any body shape as the distance to the body increases, close proximity operations require more accurate models—especially for small bodies that are highly

non-spherical.

There are two general approaches for estimating the gravity field of a body: *direct methods*, which model the gravity field directly from measurements, neglecting the physical mass distribution of the body, and *indirect methods*, which first estimate the body’s mass distribution (or more commonly, shape and density distribution), and then compute the gravity field from this physical estimate. Spherical harmonic expansions are by far the most common (direct) method used for almost all satellite navigation, which are defined by a series of harmonic coefficients. As the number of coefficients increase (i.e. their “degree” and “order”), this model converges to the true gravity field of any body with arbitrary mass distribution. However, it is only valid outside the “Brillouin sphere” (i.e. the circumscribing sphere) of the body, rendering it useless for computing the gravity at a point on the surface. Other harmonic forms exist that are valid on the body’s surface such as interior spherical harmonics [26] and the interior spherical Bessel gravity field [27], which have limited regions of convergence or are prone to large acceleration errors.

Alternatively, indirect methods for gravity field estimation attempt to construct an internal density model of the body from orbit determination (OD) data. On the one hand, this approach leverages a shape model of the body’s surface which contains much more information than harmonic coefficients alone, but on the other hand, the internal density distribution is inherently infinite dimensional and unobservable. In other words, any number of gravity measurements cannot uniquely determine the density distribution. A common approach is to start with a constant density polyhedron gravity model [28] as a first-order estimate and then augment this model with a best-fit solution of some parametric hypothesis class of density distributions [29–31]. Figure 3 shows three example density parametrizations of Comet 67P/C-G.

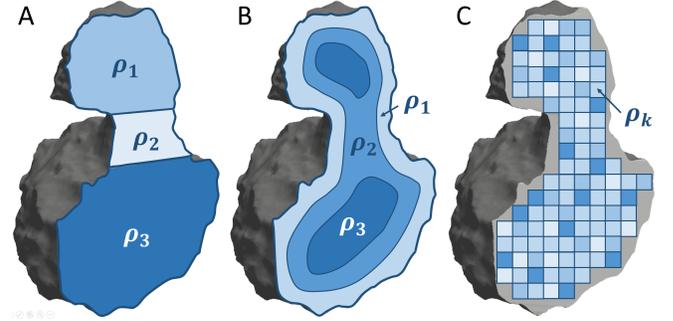


Figure 3. Examples of three hypotheses for the internal density morphology of a small body. (A) Planar partitioning into constant density sections is a common hypothesis for bodies with a suspected aggregate structure such as “contact binary” asteroids. (B) A layered structure may be assumed for bodies with a dense core or a porous surface layer. (C) A voxel decomposition offers a rich representation for arbitrary density distributions but may be severely underdetermined [31].

Density maps can either be defined based on geophysical intuition (e.g. Figs. 3A and B), or in a more general, albeit higher dimensional decomposition (Fig. 3C). For a given map, the density values of each sub-volume are determined in a least-squares sense by matching their cumulative harmonic coefficients with those measured through OD. Additionally, measuring the rotation state of the body can help to impose constraints on the density values through the inferred inertia tensor and center of mass. However, depending on the spacecraft trajectory and the quality of the OD solution, measured harmonic coefficients may only contain enough information to uniquely define a very low dimensional density

³If B is a secondary body (i.e. $N\mathbf{a}^{B_{\text{cm}}} \neq 0$), an additional term is required.

parametrization (e.g. Figs. 3A and B).

The problem with density models that are constrained to some geophysical structure is that it becomes very difficult, if not impossible to assess their accuracy—a critically important capability for reliable localization. Some error metrics have been proposed for quantifying the “goodness of fit” for such low-dimensional density models such as the “error variance,” which represents the percent error of the spherical harmonic coefficients produced by the model vs. the OD solution, and the “RMS density error,” which simply assesses the precision of a given density estimate with respect to higher resolution density estimates [31]. However, neither of these metrics are able to assess the accuracy of gravity on the surface.

On the other hand, the accuracy of generalized models such as the block decomposition in Fig. 3C is determined by the size of the tessellation and their joint covariance [31]. Takahashi et. al. analyzed the accuracy of a block-decomposed density distribution of the bi-lobed asteroid Castalia. From 90 internal blocks (fit to degree-and-order 20 harmonic coefficients), the surface gravity errors were less than 1% across most of the surface, with some regions near the “neck” of the body as high as 3% [31]. Generally, points on the surface further within the Brillouin sphere are prone to larger gravity errors.

Density modeling is a helpful intermediary for estimating surface gravity, but ultimately the accuracy is largely driven by the quality of the trajectory data. In general, more information is contained in measurements closer to the surface; distant orbits can resolve only low order harmonics, low-altitude flybys—especially within the Brillouin sphere—can constrain higher order coefficients, and low-altitude hovers or touch-and-go (TAG) maneuvers are even better. But a hopping rover equipped with a gravimeter introduces yet another extremely valuable source of information. In addition to the OD of its suborbital hopping trajectories, localized point gravity vector measurements at various locations on the surface could potentially be used to construct much higher resolution internal density maps. The formalization of this type of estimation is left for future work, but for lander/rover missions of interest in this paper, it is not unreasonable to assume an accurate surface gravity model on the order of 1%. For the case study in Sect. 6, we assume a spherical covariance, $\Sigma_g(x) = \sigma^2(x)I$, $\sigma(x) = \sigma_{\max} - |x|/R_B(\sigma_{\max} - \sigma_{\min})$, where R_B is the radius of the Brillouin sphere.

Shape Uncertainty—In addition to gravity uncertainty associated with the inaccuracies of density maps, errors in the *shape model* of the body may also contribute to gravity errors. We will consider two possible sources of errors: (1) Gaussian uncertainty in the positions of the vertices in the triangular mesh model, and (2) unmodeled features on the surface that are smaller than the resolution of the mesh. Shape reconstruction methods such as stereophotoclinometry and stereophotogrammetry, which define vertices of the triangular mesh from image data, can also compute error covariances for each vertex. This was done for the 3-million facet model (*shap5* [32]) of Comet 67P/C-G in [32], where vertex errors were found to be generally less than one or two meters on the well-observed northern hemisphere of the body with some patches of vertices as high as 5-10 meters in the southern hemisphere. The position covariance (Σ_x) of a point on the surface (i.e. on a facet) can either be interpolated exactly from the covariances of the three neighboring vertices or conservatively approximated as the largest of its neighbors’ covariances. Finally, the gravitational uncertainty associated

with the positional uncertainty can be computed as

$$\Sigma_{g_x}(x) = G(x)\Sigma_x(x)G(x)^T, \quad (10)$$

where the gravity gradient matrix, $G(x) = \nabla\nabla U(x)$, can be computed from the augmented polyhedral gravity model as in [28].

Another important consideration for characterizing gravity uncertainty on the surface are the unmodeled surface irregularities, which would generally be less than the length scale of one facet. Due to the $1/r^2$ nature of gravity dissipation, small mass fluctuations that are very close to a rover may also yield an appreciable gravity perturbation. As a rough approximation, we derive the gravity perturbation distribution produced by a field of randomly distributed rocks on a flat plane (see Fig. 4).

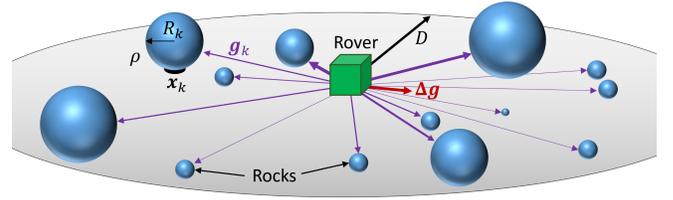


Figure 4. A statistical rock distribution is defined in the vicinity of the rover in order to estimate gravity perturbations associated with unmodeled surface features.

N spherical rocks with constant density ρ and radius distribution $f_R = \frac{\alpha(R_{\min}R_{\max})^\alpha}{R^{\alpha+1}(R_{\max}^\alpha - R_{\min}^\alpha)}$ (ranging from R_{\min} to R_{\max}) are distributed randomly over the surface disk of radius D around the rover, where $\alpha > 0$ is the power law index, which has been determined empirically for small bodies to range from 2 to 4 [33]. The gravity contribution of a single rock is distributed as $f_{|g_k|} = \frac{\beta}{|g_k|^2}$, for $\frac{4\pi G\rho R_{\min}^3}{3D^2} < |g_k| < \frac{4}{3}\pi G\rho R_{\max}$, where β is a normalizing constant. For a rock number density defined as $\gamma = \frac{N}{\pi D^2}$ (e.g. the number of rocks between R_{\min} and R_{\max} per square meter) and in the limit as $D \rightarrow \infty$, the central limit theorem⁴ can be used to determine the distribution of the net gravity perturbation (in the plane) as a Gaussian with zero mean and covariance,

$$\Sigma_{\Delta g} = \sigma_{\Delta g}^2 I_2, \quad \sigma_{\Delta g}^2 = \pi\gamma \left(\frac{4}{3}G\pi\rho\right)^2 R_{\min}^3 R_{\max}. \quad (11)$$

The magnitude of this perturbation depends on the resolution of the shape model (which influences the maximum rock size, R_{\max} , not captured in the mesh), the rock density ρ , and the number density, γ . For example, the 3-million facet model of Comet 67P/C-G (*shap5* [32]) has a mean facet area of about 15 m², so if we assume $R_{\max} = 3$ m, $R_{\min} = 0.1$ m, $\gamma = 1$ rock/m², and $\rho = 1000$ kg/m³ (roughly twice the mean density of the comet), the gravity perturbation is $\sigma_{\Delta g} \approx 3e-8$ m/s² or roughly 4 orders of magnitude lower than the mean surface gravity of 67P. Thus, in general, gravity perturbations from local surface irregularities are far less than errors induced by the gravity model.

⁴The use of the CLT assumes that g_k are independent, which is not strictly true (since, e.g. rocks cannot intersect). However, it is a reasonable assumption for sparse rock fields.

D. Total Measurement Uncertainty

In summary, the measurement model (3) is composed of the sensor model, which has a Gaussian structure with mean (4) and covariance (7), and the acceleration model (2), which is the linear combination of inertial acceleration (Gaussian with mean (8) and covariance (9)) and gravity. The gravity model has a mean, $\mu_G = \mathbf{g}(x)$, given by the gravity model of choice, and its associated covariance, Σ_G , can be augmented with shape and terrain uncertainty given by (10) and (11). Since the sensor model is an unbiased estimator of the the local acceleration (i.e. $\mu_{\bar{\mathbf{a}}|\mathbf{a}=\mathbf{a}'} = \mathbf{a}'$), the total measurement uncertainty is simply a convolution of the sensor and acceleration models: $\bar{\mathbf{a}}(x) \sim \mathcal{N}(\mu_{\bar{\mathbf{a}}}(x), \Sigma_{\bar{\mathbf{a}}}(x) + \Sigma_{\mathbf{a}}(x))$.

5. PROBABILISTIC DYNAMICS MODEL

In order to constrain the probable landing sites of a lander or rover, equation (1) suggests that the measurement model is only half of the answer. Equally important is a *prior*—that is, a belief distribution over the surface derived from prior information. In many robot localization problems, it is common to estimate the robot’s current state from an estimate of a previous state and a model of its dynamics—the most famous of which being the Kalman filter and variants thereof. However, the dynamics of a non-spherical rover bouncing on an irregular surface with largely unknown physical properties is extremely difficult (if not impossible) to model explicitly. While the dynamics of a single unperturbed ballistic trajectory are smooth and rather predictable, the surface impact dynamics are not, and moreover, when this uncertainty is compounded across a *series* of unpredictable bounces (each one highly dependent on the last), the dynamics become chaotic.

A. Uncertainty Modeling

Uncertainty in hopping dynamics can be broadly divided into three categories: (1) uncertainty in the initial state (i.e. position and velocity), (2) uncertainty in the forces acting on the rover in flight (e.g. gravity, solar radiation pressure, etc.), and (3) uncertainty in the contact forces. Ignoring bouncing dynamics, [34] proposed an analytical method for estimating the impact distribution of a single trajectory by propagating Gaussian uncertainty in the initial state and gravity model and projecting it onto a surface mesh. However, uncertainties associated with contact dynamics are particularly challenging to model because they compound uncertainties in the contact pose of the lander, the mechanical properties of the surface, and the shape of the surface itself. Through high-fidelity rigid body simulations, Van Wal, et. al. have studied how lander “settling dispersion” is affected by parameters such the lander shape [35] and surface properties [36]. Alternatively, some efforts have been made to model the contact physics on an even more detailed level through discrete element simulations of impacting regolith [37], and on the other extreme, some studies simply model the lander as a bouncing particle [34].

For the purposes of characterizing lander settling distributions, the ultimate goal of contact modeling is to accurately estimate the rebound distribution,

$$\mathbf{v}_{t+} \sim P_{\zeta}(\mathbf{v}_{t-}, \boldsymbol{\theta}_t, \zeta(\mathbf{x}_t)), \quad (12)$$

where the rebound velocity, \mathbf{v}_{t+} (optionally including rotation), depends on the impact velocity, \mathbf{v}_{t-} , the impact pose, $\boldsymbol{\theta}_t$, and a parametric description of the surface at the impact location, $\zeta(\mathbf{x}_t)$ (e.g. surface friction, rock density, regolith depth, etc.). By viewing rebounds in this statistical way,

multi-impact hopping simulations can be run *much* faster, by obviating the need to numerically integrate the contact dynamics at each impact. Of course, P_{ζ} may be highly complex and require thousands of high-fidelity simulations to estimate accurately, but for applications where many repeated landing distributions need to be computed, such as motion planning, it is far more efficient to invest the upfront effort to characterize P_{ζ} . Note, however, that a statistical rebound model does not capture the last phases of motion where persistent contact such as rolling or sliding may occur, but for non-spherical lander shapes (e.g. Philae and MASCOT), the scale of this motion is expected to be negligible compared to the distances traveled during ballistic flight.

In any parametric contact model such as Eq. (12), it is crucial that the parameters, ζ , be chosen carefully, as they can heavily influence the resulting landing distribution. For example, if the surface damping properties are lower than assumed, the actual landing distribution would likely be larger than predicted. However it is often difficult to estimate this parametric uncertainty when very little data is available, so it is better to err on the “conservative” side by biasing parameter values in a way that yields more dispersion and uniformity (e.g. higher friction and restitution). In other words, any landing distribution derived from simulations is inevitably wrong, but it is preferable to assume a model with larger variance than one with large bias. A more rigorous discussion of parametric uncertainty models is left for future work.

B. Monte Carlo Simulations of the Philae Lander

As a notional case study, we simulate the deployment of Philae from Rosetta in order to formulate a prior for localization estimates⁵. In this case, we assume that the only information available is the deployment state and uncertainty distribution, a shape model of the body, an estimate of the gravity field, the shape, mass, and mechanical properties of the lander, and optionally, information recorded by the lander during landing (e.g., the energy dissipation of the first impact measured in the landing legs, the timing of each collision, and the impulse of first touchdown as integrated from IMU data).

Based on the “500m landing ellipse” about the target landing site (assumed to be 2σ), an appropriate Gaussian distribution with spherical covariance (λI) is assumed on the deployment velocity (given by the Philae team in [38]). A reduced-order, constant-density polyhedral gravity model (10,000 facet) is assumed and the lander, modeled as a particle, experiences instantaneous impacts with a 3-million facet shape model of the comet (*shap5* [32]). Particle rebound distributions are modeled as highly stochastic: the first impact, which occurred in an upright orientation with some energy dissipation in the landing legs, was modeled as having a mean restitution of 0.4 with standard deviation of 0.1 (actual value was 0.32), and uniform angular dispersion. Subsequent collisions were assigned a mean restitution of 0.7 with a standard deviation of 0.15, and an angular dispersion of 15° (1σ about the nominal reflected velocity), which is in addition to a 1σ variability of 10° on the surface normal vector to account for unmodeled rocks and surface irregularities. Finally, the simulation is stopped probabilistically when the impact speed drops below 10 cm/s.

Figure 5 shows an example of 10,000 Monte Carlo simulations. Interestingly, even with the reduced restitution of the first impact, the surface dispersion is quite large. Some particles

⁵Note that this is a notional simulation study predicated on partial, publicly available data and not meant to represent Philae’s actual landing distribution.

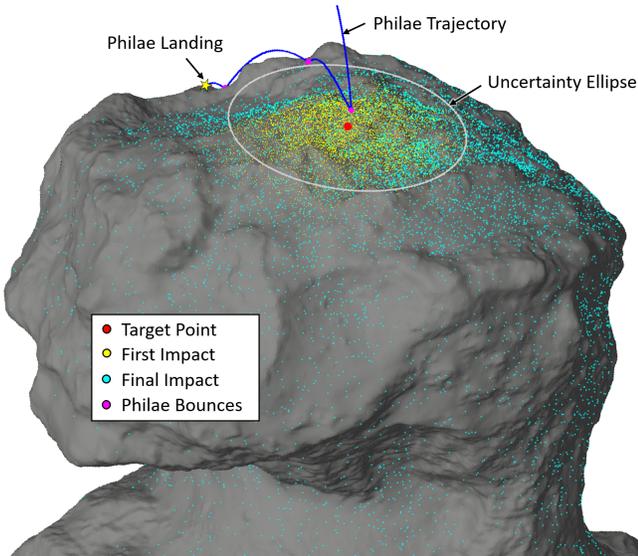


Figure 5. Monte Carlo simulations of the Philae lander deployed on to the surface of Comet 67P/C-G, where yellow dots mark the point of first touchdown and cyan dots, the final resting location. The blue line represents Philae’s actual trajectory (reconstructed by [38]) and the white circle outlines the 500 m touchdown uncertainty (assumed 2σ).

landed in the neck region of the comet, some on the complete opposite side, and 5-10% of simulations even enter orbit or escape the body entirely. This is due to the fact that the impact location has a high relative geopotential, and the lander impacted at 101 cm/s, which is above the nominal local escape velocity of 83 cm/s (which actually varies from 50 cm/s to 115 cm/s depending on the rebound direction).

C. Estimating Landing Distributions

While scatter plots such as Fig. 5 provide a visual description of the landing distribution, we need a functional form of the probability density that can be evaluated at any arbitrary point on the surface and thus, provide a quantitative prior for localization estimates.

Density estimation techniques fall into two categories: parametric and non-parametric methods. Parametric methods assume some underlying structural form of the density, such as a mixture of Gaussians, and they are described by a parameter vector, $\theta \in \mathbb{R}^k$, which is often fit to data using maximum likelihood techniques. On the other hand, non-parametric methods make no structural assumptions on the underlying density, but rather define it *implicitly* through the data itself (e.g. through a sum of kernels). Parametric models are often much cheaper to evaluate but rely on some intuition or knowledge about the underlying structure, such as the number of Gaussian components. Landing distributions are often highly irregular and thus better suited for non-parametric density estimation.

As an visual example, a multivariate kernel density model was used to evaluate the pdf at the center of each facet in the surface mesh based solely on the (x, y, z) settling position⁶ of 100,000 simulated landers (see Fig. 6). A standard multivariate normal kernel was used with a uniform, spherical bandwidth selection (i.e. $\mathbf{H} = s^2\mathbf{I}$) and a pseudo-optimal bandwidth selection

⁶Two-dimensional surface parameterizations are also possible, but they are prone to singularities and require complex map projections for irregular manifolds such as the surface of Comet 67P/C-G.

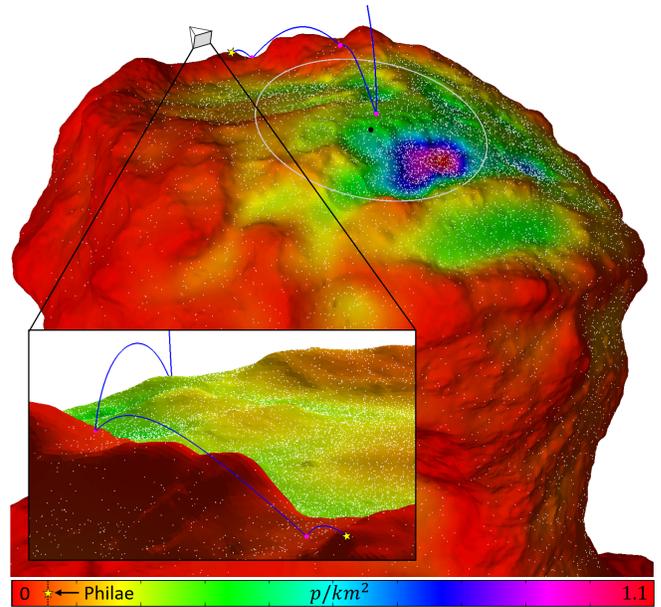


Figure 6. Multivariate kernel density estimate of the landing distribution of Philae projected onto the 3-million facet surface model of Comet 67P/C-G. Spherical Gaussian kernels are used with pseudo-optimal bandwidth of 60 m.

of $s = 60$ m (based on Silverman’s rule of thumb and cross validation). Overall, the model is an excellent fit, but adaptive kernel sizing would help in regions with sparse data. Not surprisingly, the settling distribution in Fig. 6 indicates higher probabilities in the vicinity of the touchdown site and within local gravity wells. Interestingly however, it also suggests that Philae’s actual settling location was rather unlikely. It may be tempting to blame this on a bad model, but we arrive at the same conclusion for a wide variety of contact model parametrizations.

Model accuracies notwithstanding, the distribution in Fig. 6 is solely derived from the final location of the simulated data, ignoring any intermediate information. For example, an onboard IMU should easily be able to provide fairly accurate timing and possibly acceleration data for each collision. Therefore, by also labeling simulated trajectories with this information, we can derive a *conditional* density estimate. For kernel-based densities, this can be achieved either through importance weighting or by using higher dimensional kernels. Figure 7 shows two example conditional distributions for the Philae landing: one conditioned on the timestamps of the second and final impacts and one conditioned on the impulse vector of the first impact. Reasonable bandwidth values for time and impulse were chosen based on the size of the simulated dataset and expected measurement uncertainties (30 minutes and 5 cm/s, respectively).

The time-conditioned distribution (Fig. 7B) effectively constrains the *distance* of likely settling locations from the touchdown point. In this instance, Philae traveled quite far—about 1.3 km in 2 hours—so while more distant locations such as the neck region of the comet are far less likely when conditioned on time, the distribution largely remained unchanged. On the other hand, conditioning on the impulse of the first touchdown (Fig. 7C) imposes a much tighter constraint on the landing site, as the first rebound has the largest effect on the general direction of motion.

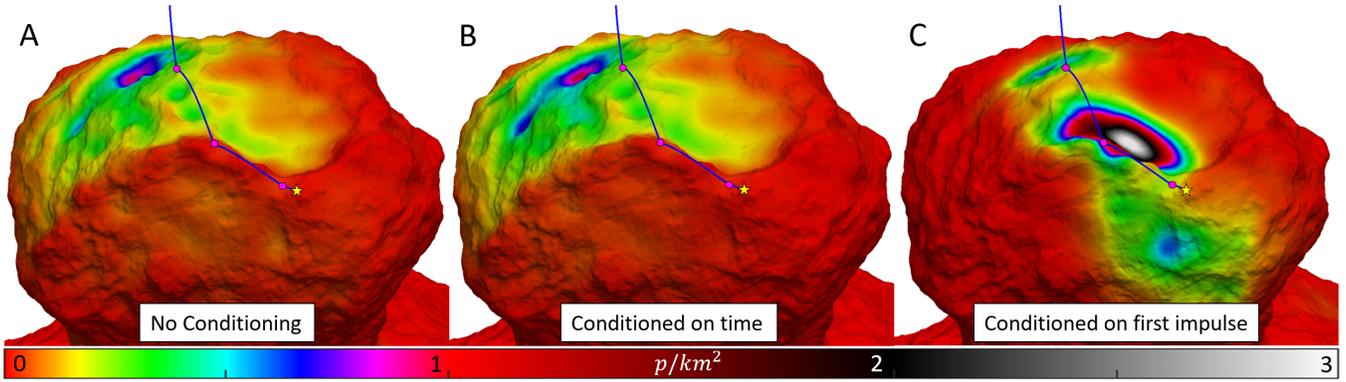


Figure 7. Conditional density estimates of the Philae landing site projected onto the 3-million facet surface model of Comet 67P/C-G. (A) The original distribution from Fig. 6 without conditioning. (B) Settling distribution conditioned on the timestamps of Philae’s second and final impacts. (C) Settling distribution conditioned on the surface impulse of Philae’s first touchdown.

6. CASE STUDY: THE PHILAE LANDER

To evaluate the potential efficacy of gravimetric localization, we study the case of the Philae lander, which, as discussed in Sect. 1, experienced a landing mishap, whereby the lander bounced over 1 km off target and was unequipped to localize itself on the surface. The question we will address here is, had Philae been equipped with a gravimeter (e.g. VEGA [25]), would it have been able to localize itself without any support from Rosetta or mission control? Of course, a gravimeter alone will never provide the localization precision of an “eye-in-the-sky” mothership (or potentially an on-board vision system), but given the information we know about the comet, sensors, and lander dynamics, to what degree could a gravimeter have constrained its location?

There are two critical elements to evaluating the localization performance: (1) what is the expected sum total of all acceleration covariances in the measurement model, and (2) how much does the nominal gravity signal vary spatially across the surface. For (1), we attempt to estimate the 1σ acceleration errors associated with each parameter. Table 1 summarizes the assumed raw error values and their computed effect on the acceleration covariance, $\Sigma_{\bar{a}}$. These values are

Source	Parameter	Reference Value	Notes	Net accel. error (ng)
Body	B_{cm}	10 m	Normal to spin axis	20
	ξ_z	0.01°	Corresponds to star tracker timestamp error of 100ms	2
	ξ_x, ξ_y	0.1°	Spin axis orientation	25
	ω_B	$10^{-7} s^{-1}$	Corresponds to period timing error of 30s	8
	α_B	$10^{-13} s^{-2}$	Estimated nutation	< 0.1
	x	5 m	Shape model vertex error	10
Star Tracker	ϵ_z	25 arcsec	Boresight axis	2
	ϵ_x, ϵ_y	2.5 arcsec	Cross-boresight axis	0.2
Gravimeter	\bar{a}_i	5 ng	VEGA	5
Gravity Model	g	2%	Percent w.r.t. the true gravity	300
	g_x	$\sigma_x = 5 m$	error along surface normal	60
	Δg	see Sect. 4C	Rock model in Sect. 4C	5

Table 1. Estimated error values (1σ) for all parameters in the measurement model and their calculated effect on the acceleration covariance, $\Sigma_{\bar{a}}$.

estimated for Comet 67P/C-G, the NST-1 nano star tracker,

the VEGA gravimeter, and where applicable, using the actual position of the Philae lander [25, 38]. Note however, that while these values are best estimates of true errors, they should be regarded as *reference* values rather than absolute truth. In other words, another scenario with different error values would yield different acceleration uncertainty accordingly.

The right column of Table 1 (shaded by value) indicates what are the largest contributors to the acceleration uncertainty, and by far, the leading contributor is uncertainty in the gravity model⁷. Errors from the sensors and the body are relatively low, and since the total acceleration error is the sum of the squares of these values (i.e. summed covariances), they are negligible compared to the gravity model.

With an understanding of the uncertainty in the measurement model, the next question is, what is the size of the corresponding uncertainty region on the body? Without considering any dynamics prior, Fig. 8 illustrates gravity uncertainty contours as a percentage of the nominal model at a reference point x^* near the Philae landing site. Encouragingly, error regions within 10% are quite localized about the reference point, with a 5% error radius of only 15 m, or roughly 0.001% of the comet’s surface area. This suggests that gravity error models within a few percent may be sufficient for high quality localization estimates.

Finally, we use the measurement model derived in Sect. 4 with values from Table 1 and the fitted prior distribution derived in Sect. 5 (without conditioning) to compute the total localization likelihood map (Eq. (1)) of the Philae lander. Figure 9 shows the computed posterior distribution projected onto the surface as a colormap. The associated 2σ gravimetry ellipse is shown along with the search ellipse derived from radar signals received from Philae’s CONSERT instrument [38]. Also shown (in green) is the preliminary search region, where follow-up images were taken by Rosetta. Interestingly, not only does the gravimetry ellipse contain Philae’s actual landing site, but it also imposes a much tighter constraint on the search space than the CONSERT ellipse alone, suggesting that an onboard gravimeter may have aided in the quest to image the Philae lander.

⁷While the 2% gravity error is not based explicitly on any specific gravity model, it is an estimate based on the modeling results in [31], which is likely conservative since Philae’s landing site is quite close to the Brillouin sphere.

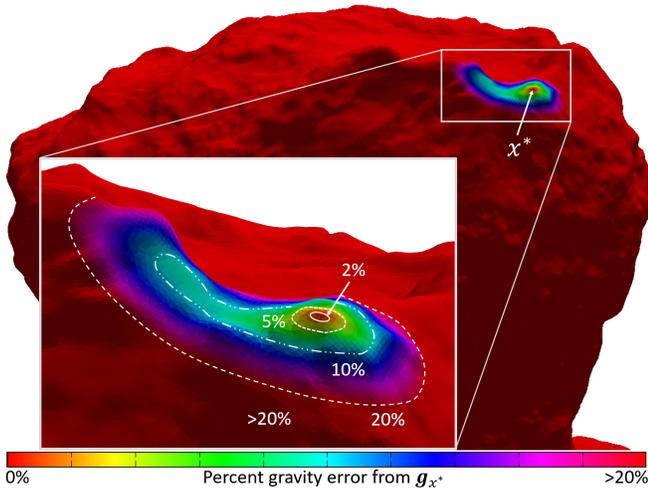


Figure 8. Spatial gravity fluctuations about a reference value of g_{x^*} at x^* are shown as a colormap with contours corresponding to 2%, 5%, 10%, and 20% deviations.

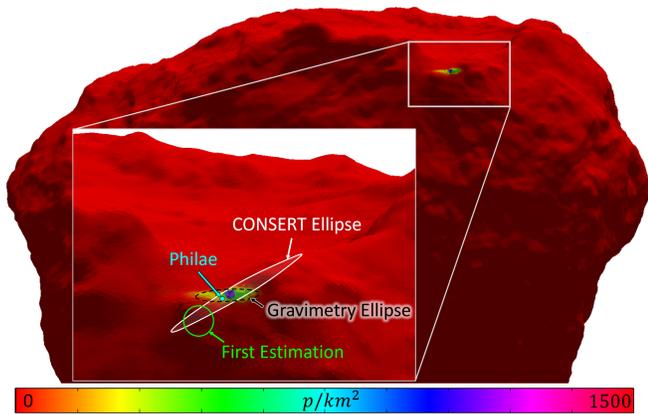


Figure 9. The posterior likelihood map of the Philae lander’s final resting location using only a gravimeter and star tracker, which is derived from error models summarized in Table 1. The uncertainty ellipse from the CONSERT instrument (white) and the preliminary search region (green) are also shown [38].

7. CONCLUSIONS

In this paper, we presented a novel localization architecture for landers and rovers on small bodies based on the use of an on-board gravimeter and star tracker. We cast the “single-shot” localization as a Bayesian estimation problem (1), where the measurement model (3) accounts for uncertainties in the sensors (7), body parameters (9), and gravity model (10), (11). We then proposed a kernel-based density estimation method for fitting a prior distribution to simulation data. Finally the Bayesian estimation framework was applied to the case of the Philae lander, suggesting that a gravimeter can produce quite localized likelihood estimates even for gravity models with moderate uncertainty.

This work leaves numerous open questions for future work. First and foremost, a formal framework is required for refining gravity and internal density models given a set of localized surface gravity measurements (e.g. leveraging gravity inversion methods from terrestrial geophysics for small bodies). Also, the use of a gravity gradiometer—an instrument that can measure components of the local gravity gradient tensor—may enhance localization capabilities and should be investigated

further. There are also other implicit pieces of information given by the gravimeter and star tracker that may help refine localization estimates, such as information about the local surface slope, or the expected star field visibility given the body pose and shape model. Another interesting piece of information that may be provided by the gravimeter is the time-varying component of the local acceleration, which may be useful on bodies with complex rotation.

The case study of the Philae lander discussed here is just a preliminary feasibility study for gravimetric localization. It should also be studied in the context of hopping rovers and on bodies of varying size and rotation. Finally, it is unlikely that a gravimeter would ever be the sole instrument used for localization, so it is also important to study the interaction and potential synergies between this method and more traditional vision-based approaches.

REFERENCES

- [1] National Research Council. Decadal survey vision and voyages for planetary science in the decade 2013–2022. Technical report, National Academy Press, 2011.
- [2] National Research Council. Defending planet earth: Near-earth-object surveys and hazard mitigation strategies. Technical report, National Academy Press, 2010.
- [3] M. J. Sonter. The technical and economic feasibility of mining the near-earth asteroids. *Acta Astronautica*, 41(4-10):637–647, 1997.
- [4] A. F. J. Abercromby, M. L. Gernhardt, S. P. Chappell, D. E. Lee, and A. S. Howe. Human exploration of phobos. In *IEEE Aerospace Conference*, 2015.
- [5] A. F. Cheng. Near-earth asteroid rendezvous: Mission summary. In *Asteroids III*. The University of Arizona Press, 2002.
- [6] D. E. Brownlee, P. Tsou, J. D. Anderson, et al. Stardust: Comet and interstellar dust sample return mission. *Journal of Geophysical Research*, 108(E10), 2003.
- [7] J. Kawaguchi, A. Fujiwara, and T. Uesugi. Hayabusa—its technology and science accomplishment summary and hayabusa-2. *Acta Astronautica*, 62(10–11):639–647, 2008.
- [8] K.-H. Glassmeier, H. Boehnhardt, E. Koschny, D. and Kührt, and I. Richter. The rosetta mission: flying towards the origin of the solar system. *Space Science Reviews*, 128(1):1–21, 2007.
- [9] M. F. A’Hearn, M. J. S. Belton, W. A. Delamere, et al. Deep impact: Excavating comet tempel 1. *Science*, 310(5746):258–264, 2005.
- [10] C. T. Russell, F. Capaccioni, A. Coradini, et al. Dawn mission to vesta and ceres. *Earth, Moon, and Planets*, 101(1–2):65–91, 2007.
- [11] D. S. Lauretta and OSIRIS-REx team. An overview of the OSIRIS-REx asteroid sample return mission. In *Lunar and Planetary Science Conference*, 2012.
- [12] H. F. Levison, C. Olkin, K. S. Noll, S. Marchi, and Lucy Team. Lucy: Surveying the diversity of the trojan asteroids: The fossils of planet formation. In *Lunar and Planetary Science Conference*, 2017.
- [13] P. Lord, S. Tilley, D. Y. Oh, et al. Psyche: Journey to a metal world. In *IEEE Aerospace Conference*, 2017.
- [14] K. Kuramoto, Y. Kawakatsu, M. Fujimoto, and MMX

Study Team. Martian moons exploration (MMX) conceptual study results. In *Lunar and Planetary Science Conference*, 2017.

- [15] J. Castillo, M. Pavone, I. Nenas, and J. A. Hoffman. Expected science return of spatially-extended in-situ exploration at small Solar System bodies. In *IEEE Aerospace Conference*, 2012.
- [16] D. Mège, J. Gurgurewicz, J. Grygorczuk, Ł. Wiśniewski, and G. Thornell. The highland terrain hopper (HOPTER): Concept and use cases of a new locomotion system for the exploration of low gravity solar system bodies. *Acta Astronautica*, 121(April–May):200–220, 2016.
- [17] B. Hockman, A. Frick, R. G. Reid, I. A. D. Nenas, and M. Pavone. Design, control, and experimentation of internally-actuated rovers for the exploration of low-gravity planetary bodies. *Journal of Field Robotics*, 34(1):5–24, 2016.
- [18] Shyam Bhaskaran, Sumita Nandi, Stephen Broschart, et al. Small body landings using autonomous onboard optical navigation. *The Journal of the Astronautical Sciences*, 58(3):409–427, 2011.
- [19] Andrew E Johnson, Yang Cheng, James Montgomery, et al. Real-time terrain relative navigation test results from a relevant environment for mars landing. In *Proc. AIAA GN&C Conference*, pages 2015–0851, 2015.
- [20] Clark F Olson, Larry H Matthies, John R Wright, Rongxing Li, and Kaichang Di. Visual terrain mapping for mars exploration. *Computer Vision and Image Understanding*, 105(1):73–85, 2007.
- [21] Mark Maimone, Yang Cheng, and Larry Matthies. Two years of visual odometry on the mars exploration rovers. *Journal of Field Robotics*, 24(3):169–186, 2007.
- [22] Tetsuo Yoshimitsu, Takashi Kubota, Ichiro Nakatani, Tadashi Adachi, and Hiroaki Saito. Micro-hopping robot for asteroid exploration. *Acta Astronautica*, 52(2):441–446, 2003.
- [23] Edmond Wai Yan So, Tetsuo Yoshimitsu, and Takashi Kubota. Divergent stereo visual odometry for a hopping rover on an asteroid surface. In *Proceedings of the International Symposium on Space Artificial Intelligence, Robotics, and Automation for Space (i-SAIRAS)*, 2010.
- [24] H. O. Seigel, I. Brcic, and P. Mistry. A guide to high precision land gravimeter surveys. *Scintrex Limited*, 222(1):1–122, 1995.
- [25] K. A. Carroll, H. Spencer, and R. E. Zee. An asteroid lander/rover for asteroid surface gravity surveying. In *AIAA/UTU Small Satellite Conference*, 2016.
- [26] Y. Takahashi, D. J. Scheeres, and R. A. Werner. Surface gravity fields for asteroids and comets. *AIAA Journal of Guidance, Control, and Dynamics*, 36(2):362–374, 2013.
- [27] Y. Takahashi and D. J. Scheeres. Small body surface gravity fields via spherical harmonic expansions. *Celestial Mechanics and Dynamical Astronomy*, 119(2):169–206, 2014.
- [28] R. A. Werner and D. J. Scheeres. Exterior gravitation of a polyhedron derived and compared with harmonic and mascon gravitation representations of asteroid 4769 castalia. *Celestial Mechanics and Dynamical Astronomy*, 65(3):313–344, 1996.
- [29] R. S. Park, R. A. Werner, and S. Bhaskaran. Estimating small-body gravity field from shape model and navigation data. *AIAA Journal of Guidance, Control, and Dynamics*, 33(1):212–221, 2010.
- [30] R. P. Russell and N. Arora. Global point mascon models for simple, accurate, and parallel geopotential computation. *AIAA Journal of Guidance, Control, and Dynamics*, 33(5):1568–1581, 2012.
- [31] Y. Takahashi and D. J. Scheeres. Morphology driven density distribution estimation for small bodies. *Icarus*, 233(1):179–193, 2014.
- [32] L. Jorda, R. Gaskell, C. Capanna, et al. The global shape, density and rotation of comet 67p/churyumov-gerasimenko from preperihelion rosetta/osiris observations. *Icarus*, 277:257–278, 2016.
- [33] S. Tardivel, D. J. Scheeres, P. Michel, S. Van wal, and P. Sánchez. Contact motion on surface of asteroid. *AIAA Journal of Spacecraft and Rockets*, 51(6):1857–1871, 2014.
- [34] B. Hockman and M. Pavone. Stochastic motion planning for hopping rovers on small solar system bodies. In *Int. Symp. on Robotics Research*, 2017. In Press.
- [35] S. Van wal, S. Tardivel, and D. J. Scheeres. High-fidelity small body lander simulations. In *Int. Conf. on Astrodynamics Tools and Techniques*, 2016.
- [36] S. Van wal, S. Tardivel, and D. J. Scheeres. Parametric study of ballistic lander deployment to small bodies. *AIAA Journal of Spacecraft and Rockets*, 0(0):1–26, 2017.
- [37] F. Thuillet, C. Maurel, P. Michel, et al. Numerical simulations of surface package landing on a low-gravity granular surface: application to the landing of mascot onboard hayabusa 2. In *Lunar and Planetary Science Conference*, 2017.
- [38] J. Biele, S. Ulamec, M. Maibaum, et al. The landing(s) of philae and inferences about comet surface mechanical properties. *Science*, 349(6247):1–6, 2015.

BIOGRAPHY



Benjamin Hockman is a Ph.D. candidate in Mechanical Engineering at Stanford University. He received his B.Eng. in Mechanical Engineering with a concentration in Aerospace from the University of Delaware, and a M.Sc. in Mechanical Engineering from Stanford University in 2015. He is currently an NSF fellow. Ben's research interests include design, control, and motion planning of spacecraft and rovers in extreme and uncertain environments, learning-based control and navigation for robots with highly stochastic dynamics, dynamic grasping and object manipulation, and design of novel robotic systems.



Dr. Robert G. Reid is a robotics engineer in the Robotic Mobility Group at NASA's Jet Propulsion Laboratory (JPL), specializing in autonomous systems for extreme terrains, robotic perception, simultaneous localization and mapping (SLAM) and multi-robot SLAM. He has worked on a variety of autonomous systems, including driving, flying, hopping and rappelling robots. At JPL, he has

developed vision-based SLAM systems for robots that can hop and tumble over the surface of comets and asteroids, and autonomous quadrotors capable of agile flight through cluttered environments. He received a B.Eng from Curtin University, and a Ph.D. in large-scale multi-robot SLAM from the University of Western Australia.



Dr. Issa A. D. Nesnas is a principal technologist and the supervisor of the Robotic Mobility group at the Jet Propulsion Laboratory with over two decades of research in space robotics and industrial automation. He leads research in autonomy and mobility with a focus on extreme terrain access and microgravity mobility. He contributed to the development of autonomous rover navigation and visual

target tracking and participated in the development of the Curiosity and Mars 2020 rovers. Dr. Nesnas served on NASA's Capability Leadership Team for Autonomy and was the co-chair for NASA's Technology Roadmaps for Robotics and Autonomous Systems. He holds several patents and has over fifty publications in this field. He holds a B.E. degree in Electrical Engineering from Manhattan College and a M.S. and Ph.D. in Mechanical Engineering with a specialization in robotics from the University of Notre Dame.



Dr. Marco Pavone is an Assistant Professor of Aeronautics and Astronautics at Stanford University, where he is the Director of the Autonomous Systems Laboratory. Before joining Stanford, he was a Research Technologist within the Robotics Section at the NASA Jet Propulsion Laboratory. He received a Ph.D. degree in Aeronautics and Astronautics from the Massachusetts Institute of Technology in 2010.

Dr. Pavone's areas of expertise lie in the fields of controls and robotics. His main research interests are in the development of methodologies for the analysis, design, and control of autonomous systems, with an emphasis on autonomous aerospace vehicles and large-scale robotic networks. He is a recipient of a PECASE Award, an ONR YIP Award, an NSF CAREER Award, a NASA Early Career Faculty Award, a Hellman Faculty Scholar Award, and was named NASA NIAC Fellow in 2011. He is currently serving as an Associate Editor for the IEEE Control Systems Magazine.